Storage Area Networks

SANs aim to present solutions to many of the problems associated with large-scale data storage. SANs build on the server-attached model through the creation of a separate network of storage devices, independent of the organization's LAN or communications network. Storage networks can include disk drives, RAID devices, tape libraries, and other storage equipment. Multiple servers— and even client systems— can participate in the storage network to gain access to these devices. The creation of a network of storage devices offers an organization many more options than the traditional approach of connecting storage devices directly to servers.

SANs move data efficiently without adding to the load of the communications network. While local-area networks (LANs) use protocols such as Ethernet and TCP/IP, SANs currently rely on Fibre Channel technologies (though SAN over IP is another upcoming technology). In its original and simpler form, Fibre Channel-Arbitrated Loop (FC-AL) connects a limited number of devices in a physical ring topology. More sophisticated SANs use hubs, switches, and routers to create a Fibre Channel fabric that can include a very large number of devices spanning extensive geographical locations.

The key principle of SAN involves offloading data transfers from the communications network. In a SAN environment, many data operations can be accomplished without having to traverse the LAN. Data communications protocols such as Ethernet and IP introduce significant overhead in the transmission, while Fibre Channel technologies dispatch large amounts of data with great efficiency.

The SAN is transparent to network end users, so client computers and Web users are unaware of the existence of the back-end SAN architecture. From the perspective of users on the LAN, there are no outward signs to indicate that the network servers rely on a SAN for access to storage devices.

The Pros and Cons of SAN

A major limitation of the traditional server-attached storage model lies in the restrictions in the numbers of storage devices that can be added to any given server and the level of difficulty in increasing storage capacities among groups of servers. The SCSI architecture does not scale well to very large-capacity storage systems. The number of devices per SCSI chain and slots available for host bus adapters per server are often too low to achieve the overall capacities needed by many organizations. A basic FC-AL system, however, supports up to 126 devices, while Fibre Channel fabrics can support an unlimited number of devices.

SANs allow servers to be physically separate from their storage devices. While SCSI cabling requires devices to be within a few feet of each other, Fibre Channel supports distances of many kilometers. This flexibility allows the centralization of storage equipment even when servers are distributed. Through a SAN, an organization can create large vaults of centralized storage that can be associated with servers that are housed centrally or are dispersed through the organization. SANs allow organizations

with multiple data centers to distribute access to storage assets independently of physical location.

While the underlying technologies differ significantly, network managers will see many similarities between LAN and SAN architectures and equipment. Fibre Channel hubs, for example, allow a number of devices to be interconnected, but share the overall bandwidth of the unit. Fibre Channel switches deliver faster performance and offer each device dedicated bandwidth. In a large SAN, Fibre Channel switches connect to each other in redundant paths to form a physical mesh network that continues to operate even when individual links fail. These interconnected switches form a Fibre Channel fabric that allows each of the storage devices and servers to communicate with one another through redundant paths. Fibre channel bridges connect non-Fibre Channel devices, especially SCSI devices, to the SAN.

SANs can include tape libraries, which can be used to back up all the other storage devices within the fabric or loop. In a LAN-only environment, backup operations can saturate the network. SANs, however, enable backup operations to occur without moving any data on the network and can even occur without the intervention of a server. Serverless backup, or the direct copying of data from a disk system to a tape library on a SAN, is one of the major benefits of this architecture.

The downside of SANs lies in their cost and complexity. IT departments will need to include technical staff with expertise to manage the SAN in addition to those that support existing LANs and wide-area networks (WANs). The fibre optic network and the communications equipment that comprise the SAN require significant investment as well. Yet for many organizations, the overhead of implementing and managing a SAN is less onerous than the problems associated with expanding file servers with dedicated storage to keep up with the demands of an organization's large-scale storage requirements.

Another limitation of the SAN model involves the inability for multiple servers to attach to the same device for direct file sharing. With NAS (network-attached storage), file sharing is commonplace, since multiple servers— or even clients— can mount the exported file systems and access the same files. Application-level software takes care of file and record locking to prevent data corruption. New approaches are being developed to address this limitation in SAN architecture as part of the general trend toward convergence of SAN and NAS technologies.

Excerpt from “NAS vs. SAN Overview,” GartnerGroup, by April Adams and Marshall Breeding, 16 April 2001